



[Home](#) › [Statistics](#) › [Detailed methodology information](#) › [Concepts, sources and methods](#)  
› [Labour Statistics: Concepts, Sources and Methods](#) › [2023](#) › [Methods: Four pillars of labour statistics](#)

## ✓ Latest release

# Methods: Four pillars of labour statistics

## Labour Statistics: Concepts, Sources and Methods

Reference period: 2023

Released 24/11/2023

### Sections

#### ▼ [Labour Statistics: Concepts, Sources and Methods, 2023](#)

##### [Overview](#)

##### ▶ [Concepts and sources](#)

#### ▼ [Methods: Four pillars of labour statistics](#)

##### ▶ [Household surveys](#)

##### [Business surveys](#)

##### ▶ [Administrative data](#)

##### ▶ [Australian Labour Account](#)

##### [Glossary](#)

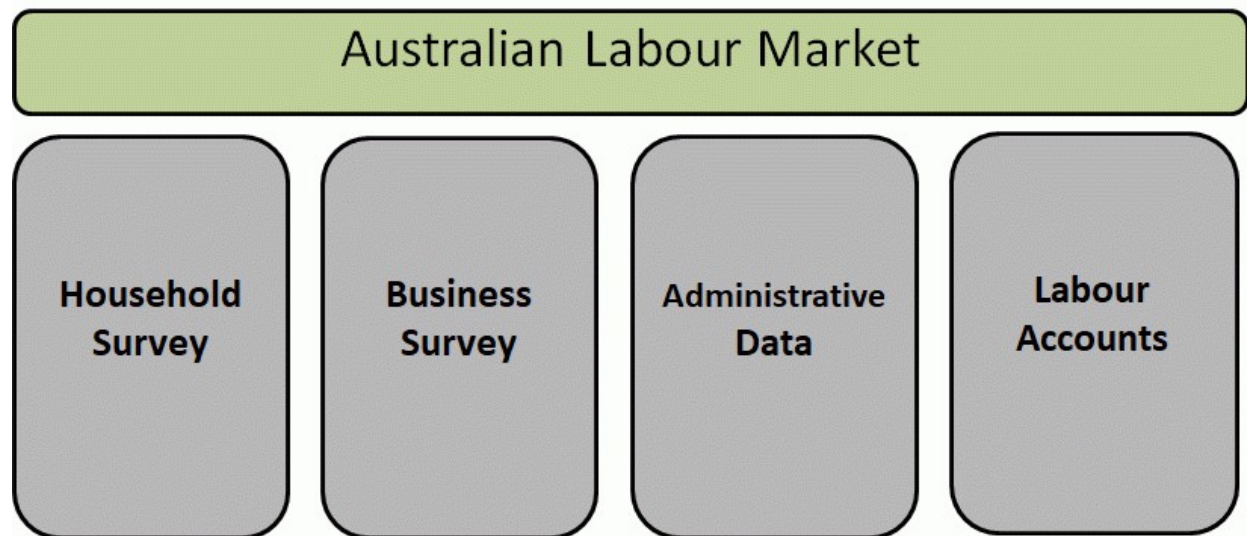
##### [Acronyms](#)

Close

ABS labour statistics are drawn from four key types of data sources, or “pillars” of data, which provide complementary insights into the labour market. These are:

- household surveys - individual households answer labour market questions about their individual, family or household circumstances (e.g. the monthly Labour Force Survey)
- business surveys - collect a broad range of information from businesses about jobs and employees (e.g. the Survey of Employee Earnings and Hours, Job Vacancies Survey)
- administrative data - information maintained by governments (such as taxation data) and other entities made available to the ABS for statistical purposes (e.g. as published in Weekly Payroll Jobs)

- accounts compilation - bringing together data from separate administrative, business, and household sources to produce an Australian Labour Account



## Sample surveys versus censuses

The ABS uses both sample surveys and censuses to collect information from a population about characteristics of interest. In the field of labour statistics, the ABS uses sample surveys of households and businesses, as well as censuses (such as the Industrial Disputes collection).

Censuses involve the collection of information from all units in the target population, while sample surveys involve the collection of information from only a part (sample) of the target population.

Sample surveys have both advantages and disadvantages when compared with censuses. Some advantages are reduced costs (as less time is needed to collect, process and produce data), possible reductions in non-sampling error (this concept is discussed in further detail later in this chapter), improved timeliness, and the potential to gather more detailed information from each respondent.

A disadvantage of sample surveys is that estimates are subject to sampling error, which occurs because data were obtained from only a sample rather than the entire population (this concept is discussed in further detail later in this chapter). Also, as a result of obtaining only a small number of observations in particular geographical areas and sub-populations, detailed cross-tabulations may be subject to high levels of error and be of limited use.

Censuses are generally used when broad level information is sought for many fine sub-groups of the population, whereas sample surveys are used to collect detailed information to estimate for broader levels of the population.

---

## Sample design and sampling techniques

ABS labour-related household and business sample surveys use probability sampling techniques, drawing their samples from a population frame. This section briefly defines and explains key concepts and terms related to survey design. See the household and business surveys sections for more detail on aspects of survey design that are particular to these types of surveys.

### Population

A survey is concerned with two types of population: the target population, and the survey population. The target population is the group of units about which information is sought, and is also known as the scope of the survey. It is the population at which the survey is aimed. The scope should state clearly the units from which data are required and the extent and time covered, e.g. households (units) in Australia (extent) in August 2020 (time).

However, the target population is a theoretical population, as there are usually a number of units in the target population which cannot be surveyed. These include units which are difficult to contact and units which are missing from the frame. The survey population is that part of the population that is able to be surveyed, and is also called the coverage population.

### Statistical units

Statistical units are used in the design, collection, analysis and dissemination of statistical data. There are several types of units, including: sampling units (the units selected in the sample survey), collection units (the units from which data are collected), reporting units (the units about which data are collected), and analysis units (the units used for analysis of the data). The units used in a survey may change at various stages in the survey cycle. For example, the Labour Force Survey uses a sample of households (sampling unit) from which information is collected from any responsible adult (collection unit) about each person in the household in scope of the survey (reporting units). The results of the survey may then be analysed for families (analysis unit).

### Frames

The frame comprises a list of statistical units (e.g. persons, households or businesses) in the population, together with auxiliary information about each unit. It serves as a basis for selecting the sample. Two types of frames are used in ABS labour-related surveys:

- List based frames - List based frames comprise a list of all sampling units in the survey population. List based frames are commonly used in surveys of businesses. ABS

business surveys currently draw their list frames from the ABS Business Register.

- Area based frames - Area based frames comprise a list of non-overlapping geographic areas. These areas may be defined by geographical features such as rivers and streets. They are usually used in household surveys. Once an area is selected, a list is made of the households in the area, and a sample of households selected from the list. Examples of geographic areas that may be used to create area frames include: local government areas; census collection districts; and postcodes.

Auxiliary variables are characteristics of each unit for which information is known on the frame prior to the survey. Auxiliary variables can be used in the sample design to better target the population of interest, if the information on the frame is of sufficiently high quality and is correlated with the variables of interest in the survey. They can also be used in the estimation process in conjunction with the survey data: for example, industry of businesses.

For most sampling methodologies, it is desirable to have a complete list from which to select a sample. However, in practice it can be difficult to compile such a complete list and therefore frame bias may be introduced. Frame bias occurs when an inappropriate frame is used or there are problems with the composition of the frame, with the result that the frame is not representative of the target population. Frames become inaccurate for many reasons. One of the most common problems is that populations change continuously, causing frames to become out of date. Frames may also be inaccurate if they are compiled from inaccurate sources. The following are some of the problems that can occur in the composition of frames.

Under coverage occurs when some units in the target population that should appear on the frame do not. These units may have different characteristics from those units which appear on the frame, and therefore results from the survey will not be representative of the target population.

Out of scope units are units that appear on the frame but are not elements of the target population. Selection of a number of out of scope units in the sample reduces the effective sample size, and increases sampling error. Furthermore, out of scope units appearing on the frame may be incorrectly accounted for in the estimation process, which may lead to bias in survey estimates.

Duplicates are units that appear more than once on the frame. The occurrence of duplicates means that the probability of selection of the units on the frame is not as it should be for the respective sample design. In particular, the duplicate units will have more than the correct chance of selection, introducing bias towards the characteristics of these units. Duplicates also increase sampling error.



Deaths are units that no longer exist in the population but are still on the frame. Deaths have the same impact on survey results as out of scope units.

The quality of auxiliary variables can affect the survey estimates of the variables of interest, through both the survey design and the estimation process.

The ABS attempts to minimise frame problems and uses standardised sample and frame maintenance procedures across collections. Some of the approaches taken are to adjust estimates using new business provisions, and to standardise across surveys the systems for handling estimation, imputation and outliers.

## **Probability samples**

Probability samples are samples drawn from populations such that every unit in the population has a known, or calculable, non-zero probability of selection which can be obtained prior to selection. In order to calculate the probability of selection, a population frame must be available. The sample is then drawn from this frame. Alternatives to probability samples are samples formed without a frame, such as phone-in polls.

Probability sampling is the preferred ABS method of conducting major surveys, especially when a population frame is available. Probability samples allow estimates of the accuracy of the survey estimates to be calculated. They are also used in ABS surveys as a means of avoiding bias in survey results. Bias is avoided when either the probability of selection is equal for all units in the target population or, where this is not the case, the effect of non-equal probabilities is allowed for in estimation.

## **Stratified sampling**

Stratified sampling is a technique which uses auxiliary information available for every unit on the frame to increase the efficiency of a sample design. Stratified sampling involves the division (stratification) of the population frame into non-overlapping, homogeneous (similar) groups called strata, which can be treated as totally separate populations. A sample is then selected independently from each of these groups, and can therefore be selected in different ways for different strata, e.g. some strata may be sampled using 'simple random sampling' while others may be 'completely enumerated'. These terms are explained below. Stratification variables may be geographical (e.g. State, capital city/balance of State) or non-geographical (e.g. number of employees, industry, turnover).

All surveys conducted by the ABS use stratification. Household surveys use mainly geographic strata. Business surveys typically use strata which are related to the economic activity undertaken by the business, for example industry and size of the business (the latter based on employment size).

## **Completely enumerated strata**

Completely enumerated strata are strata in which information is obtained from all units. Strata that are completely enumerated tend to be those where: each population unit within the stratum is likely to contribute significantly to the estimate being produced (such as strata containing large employers where the estimate being produced is employment); or there is significant variability across the population units within the stratum.

## **Simple random sampling**

Simple random sampling is a probability sampling scheme in which each possible sample of the required size has the same chance of selection. It follows that each unit of the population has an equal chance of selection.

Simple random sampling can involve units being selected either with or without replacement. Replacement sampling allows the units to be selected multiple times, whereas without replacement sampling allows a unit to be selected only once. In general, simple random sampling without replacement produces more accurate results as it does not allow sample to be 'wasted' on duplicate selections. All ABS surveys that use simple random sampling use the 'without replacement' variant. Simple random sampling without replacement is used in most ABS business surveys.

## **Systematic sampling**

Systematic sampling is used in most ABS household surveys, and provides a simple method of selecting the sample. It involves choosing a random starting point within the frame and then applying a fixed interval (referred to as the 'skip') to select members from a frame.

Information on auxiliary variables can be used in systematic sampling to improve the efficiency of the sample. The units in the frame can be ordered with respect to auxiliary variables prior to calculating the skip interval and starting point. This approach ensures that the sample is spread throughout the range of units on the frame, ensuring a more representative sample with respect to the auxiliary variable.

Systematic sampling with ordering by auxiliary variables is only useful if the frame contains auxiliary variables about each of the units in the population, and if these variables are related to the variables of interest. The relationship between the variables of interest and the auxiliary variables is often not uniform across strata. Consequently, it is possible to design a sample survey with only some of the strata making use of auxiliary variables.

## **Probability proportional to size sampling**

Probability proportional to size sampling is a selection scheme in which units in the population do not all have the same chance of selection. With this method, the larger the

unit with respect to some measure of size, the greater the probability that unit will be selected in the sample. Probability proportional to size sampling will lead to unbiased estimates, provided the different probabilities of selection are accounted for in estimation.

### **Cluster sampling**

Cluster sampling involves the units in the population being grouped into convenient clusters, usually occurring naturally. These clusters are non-overlapping, well-defined groups which usually represent geographical areas. The sample is selected by selecting a number of clusters, rather than directly selecting units. All units in a selected cluster are included in the sample.

### **Multi-stage sampling**

Multi-stage sampling is an extension of cluster sampling. It involves selecting a sample of clusters (first-stage sample), and then selecting a sample of population units within each selected cluster (second-stage sample). The sampling unit changes at each stage of selection. Any number of stages can be employed. The sampling units for any given stage of selection each form clusters of the next-stage sampling units. Units selected in the final stage of sampling are called final-stage units (or ultimate sampling units). The Survey of Employee Earnings and Hours uses multi-stage sampling - businesses (the first-stage units) selected in the survey are asked to select a sample of 'employees' (the final-stage units) using employee payrolls. Household surveys also use multi-stage sampling.

### **Multi-phase sampling**

Multi-phase sampling involves collecting basic information from a sample of population units, then taking a sub-sample of these units (the second-phase sample) to collect more detailed information. The second-phase sample is selected using the information collected in the first phase, and allows the second-phase sample to be targeted to the specific population of interest. Population totals for auxiliary variables, and values from the first-phase sample, are used to weight the second-phase sample for the estimation of population totals.

Multi-phase sampling aims to reduce sample size and the respondent burden and collection costs, while ensuring that a representative sample is still selected from the population of interest. It is often used when the population of interest is small and difficult to isolate in advance, or when detailed information is required. Multi-phase sampling is also useful when auxiliary information is not known for all of the frame units, as it enables the collection of data for auxiliary variables in the first-phase sample.

The first-phase sample is designed to be large to ensure sufficient coverage of the population of interest, but only basic information is collected. The basic information is then

used to identify those first-phase sample units which are part of the population of interest. A sample of these units is then selected for the second-phase sample. Therefore, the sampling unit remains the same for each phase of selection. If multi-phase sampling was not used, detailed information would need to be collected from all first-phase sample units to ensure reasonable survey estimates. In this way, multi-phase sampling reduces the overall respondent burden.

---

## Weighting and estimation

Sample survey data only relate to the units in the sample. Therefore, the sample estimates need to be inflated to represent the whole population of interest. Estimation is the means by which this inflation occurs.

The following section outlines various methods of calculating the population estimates from the sample survey data. It then describes various editing procedures used in labour-related statistics to improve the population estimates.

Estimation is essentially the application of weights to the individual survey, and summing these weighted records to estimate totals. The value of these weights is determined with respect to one or more of the following three factors:

- the probability of selection for each survey unit (probability weighting);
- adjustment for non-response to correct for imbalances in the characteristics of responding sample units (post-stratification); and
- adjustments to agree with known population totals for auxiliary variables - to correct for further imbalances in the characteristics of the selected sampled units (post-stratification, ratio estimation, calibration).

Weights are determined using formulae (estimators) of varying complexity.

### Number-raised estimation

Number-raised weights are given by  $N_h/n_h$  (where  $N_h$  is the total number of units in the population for the stratum, and  $n_h$  is the number of responding units in the sample for that stratum). The weight assigned to each survey unit indicates the number of units in the target population that the survey unit is meant to represent. For example, a survey unit with a weight of 100 represents 100 units in the population. Each survey unit in a stratum is given the same weight. Number-raised weights can only be used to weight simple random samples.

Advantages of number-raised estimation are: it does not require auxiliary data; it is unbiased; and the accuracy of the estimates can be calculated relatively simply. However,

number-raised estimation is not as accurate as some other methods with the same overall sample size.

## **Ratio estimation**

Ratio estimation involves the use of known population totals for auxiliary variables to improve the weighting from sample values to population estimates. It operates by comparing the survey sample estimate for an auxiliary variable with the known population total for the same variable on the frame. The ratio of the sample estimate of the auxiliary variable to its population total on the frame is used to adjust the sample estimate for the variable of interest.

The ratio weights are given by  $X/x$  (where  $X$  is the known population total for the auxiliary variable, and  $x$  is the corresponding estimate of the total based on all responding units in the sample). These weights assume that the population total for the variable of interest will be estimated by the sample equally as well (or poorly) as the population total for the auxiliary variable is estimated by the sample.

Ratio estimation can be more accurate than number-raised estimation if the auxiliary variable is highly correlated with the variable of interest. However, it is subject to bias, with the bias increasing for smaller sample sizes and where there is lower correlation between the auxiliary variable and the variable of interest.

## **Post-stratification**

Post-stratification estimation also involves the use of auxiliary information to improve the weighting from sample values to population estimates. Subgroups of the survey sample units are formed based on auxiliary variables after the survey data have been collected. Estimates of subgroup population sizes (based on probability weighting) are compared with known subgroup population sizes from independent sources. The ratio of the two population sizes for each subgroup is used to adjust the original estimate for the variable of interest (based on probability sampling).

Post-stratification is used to refine the estimation weighting process by correcting for sample imbalance and, assuming that the survey respondents are representative of missing units, correcting for non-response. For example, in the LFS, the sample is post-stratified by age, sex, capital city/rest of State, and State/Territory of usual residence. Estimates of the number of persons in these subgroups based on Census/Estimated Resident Population data are then compared to the estimates based on the survey sample to give the post-stratification weights.

## **Calibration**

Calibration essentially uses all available auxiliary information to iteratively modify the original weights (based on number-raised weights). The new weights ensure that the sample estimates are consistent with known auxiliary information. Both post-stratification and ratio estimation can be used as part of the calibration weighting process. Calibration is useful if the survey sample estimates need to match the unit totals for a number of different subgroups, or for more than one auxiliary variable. It is mostly used in Special Social Surveys. For example, the Survey of Employment and Unemployment Patterns was weighted so that the survey estimates aligned with both population estimates based on Census data and estimates of the number of persons 'employed', 'unemployed' and 'not in the labour force' from the LFS.

---

## Editing

Editing is the process of correcting data suspected of being wrong, in order to allow the production of reliable statistics. The aims of editing are:

- to ensure that outputs from the collection are mutually consistent: for example, two different methods of deriving the same value should give the same answer;
- to correct for any missing data;
- to detect major errors, which could have a significant effect on the outputs; and
- to find any unusual output values and their causes.

The purpose of editing is to correct non-sampling errors, such as those introduced by misunderstanding of questions or instructions, interviewer bias, miscoding, non-availability of data, incorrect transcription, non-response, and non-contact. Non-response occurs when all (total non-response) or part (partial non-response) of a questionnaire is not completed by the respondent. High levels of non-response can cause bias in the sample based estimates.

Editing is also used to identify outliers. The statistical term 'outlier' has several definitions, depending on the context in which it is used. Here it is used loosely to describe extreme values that are verified as being correct, but are very different from the values reported by similar units, and are expected to occur only very rarely in the population as a whole. In practice, an outlier is usually considered to be a unit that has a large effect on survey estimates of level, on estimates of movement, or on the sampling variance. This may occur because the unit is not similar to other units in the stratum - for example, if its' true employment is much greater than the frame employment. It may also occur when an extreme value is recorded for some variable from an otherwise ordinary sampling unit.

Certain types of non-response, and the presence of outliers in the sample, may be addressed using a variety of statistical techniques.

Imputation involves supplying a value for a non-responding unit, or to replace 'suspect' data. Imputation methods fall into three groups:

- the imputed value may be derived from other information supplied by the respondent;
- the imputed value may be derived from information supplied by other similar respondents in the current survey; and
- the values supplied by the respondent in previous surveys may be modified to derive a value.

The following imputation methods are used in labour-related surveys:

- Deductive imputation involves correcting a missing or erroneous value by using other information that reveals the correct answer. For example, a response of 18,000 has been given where respondents have been asked to reply in '\$000s' and where the expected range of responses is 13-21. A quick examination of other parts of the form shows that \$18,000 is very likely the amount actually spent by the respondent, so 18,000 is 'corrected' to 18.
- Central-value imputation involves replacing a missing or erroneous item with a value considered to be 'typical' of the sample or sub-sample concerned. Live respondent mean is an example of central-value imputation. This technique involves calculating the average stratum value for the data item of interest across all responding live units in the stratum, and assigning this value to all live non-responding units in the stratum.
- Hot-deck imputation is similar to central-value imputation, but takes the absolute value from a donor unit: for example, earnings per hour for a given combination of occupation, location and industry in Characteristics of Employment.
- Cold-deck imputation involves using previous survey data to amend items which fail edits. It may involve copying data from the previous survey cycle to the current cycle. One specific example of this type of imputation is Beta imputation, which involves estimating missing values by applying an imputed growth rate to the most recently reported data for these units, provided that data have been reported in either of the two previous periods.

When adjusting for outliers, a compromise is always necessary between the variability and bias associated with an estimate. There are two methods available for dealing with outliers. Historically the ABS has used the 'surprise outlier' approach for most business surveys, but over time has gradually changed to using 'winsorization'.

- Surprise outlier approach - Generally, this technique is used to deal with a selected unit which is grossly extreme for a number of variables. The approach treats each outlier as if it were the only extreme unit in the stratum population. The outlier is given a weight of one, as if it had been selected in a CE stratum. As a result of the outlier's movement to the CE stratum, the weight for units in the outlier's selection stratum has to be recalculated, as the population and sample size have effectively been reduced by one. This has the effect that the other population units which would have been represented



by the outlier are now represented by the average of the other units in the stratum. Therefore, the choice of treatments for a suspected outlier using the surprise outlier approach are either for it to represent all of the units it would normally represent, or to represent no units other than itself. It is preferable to set a maximum number of surprise outliers which can be identified in any one survey.

- Winsorization technique - This technique is a more flexible approach. Here a value is considered to be an outlier if it is greater than a predetermined cut off. The effect of the outlier on the estimates is reduced by modifying its reported value. On application of the winsorization formula, sample values greater than the cut off are replaced by the cut off plus a small additional amount. The additional amount is the difference between the sample value and the cut off, multiplied by the stratum sampling fraction. Thus winsorization has most impact in strata with low sampling fractions, and the impact decreases as sampling fractions increase. Effectively, winsorization results in the outlier only representing itself, with the remaining population units that would have been represented by the outlier being instead represented by the cut off.

---

## Time series estimates

Time series are statistical records of various activities measured at regular intervals of time, over relatively long periods. Data collected in irregular surveys do not form time series. The following section outlines the various elements of time series, and describes the ABS method of calculating seasonally adjusted and trend estimates.

ABS time series statistics are published in three forms: original, seasonally adjusted and trend.

Original estimates are the actual estimates the ABS derives from the survey data or other non-survey sources. Original estimates are comprised of trend behaviour, systematic calendar related influences, and irregular influences.

Systematic calendar related influences operate in a sustained and systematic manner that is calendar related. The two most common of these influences are seasonal influences and trading day influences.

Seasonal influences occur for a variety of reasons:

- They may simply be related to the seasons and related weather conditions, such as warmth in summer and cold in winter. Weather conditions that are out of character for a particular season, such as snow in summer, would appear as irregular, not seasonal, influences.
- They may reflect traditional behaviour associated with various social events (e.g. Christmas and the associated holiday season).
- They may reflect the effects of administrative procedures (e.g. quarterly provisional tax

payments and end of financial year activity).

Trading day influences refer to activity associated with the number and types of days in a particular month, as different days of the week often have different levels of activity. For instance, a calendar month typically comprises four weeks (28 days) plus an extra two or three days. If these extra days are associated with high activity, then activity for the month overall will tend to be higher.

Seasonal and trading day factors are estimates of the effect that the main systematic calendar related influences have on ABS time series. These evolve to reflect changes in seasonal and trading patterns of activity over the life of the time series, and are used to remove the effect of seasonal and trading day influences from the original estimates.

Seasonally adjusted estimates are derived by removing the systematic calendar related influences from the original estimates. Seasonally adjusted estimates capture trend behaviour, but still contain irregular influences that can mask the underlying month to month or quarter to quarter movement in a series. Seasonally adjusted estimates by themselves are only relevant for sub-annual collections.

Irregular influences are short term fluctuations which are unpredictable, and hence are not systematic or calendar related. Examples of irregular influences are those caused by one-off effects such as major industrial disputes or abnormal weather patterns. Sampling and non-sampling errors that behave in an irregular or erratic fashion with no noticeable systematic pattern are also irregular influences.

Trend estimates are derived by removing irregular influences from the seasonally adjusted estimates. As they do not include systematic, calendar related influences or irregular influences, trend estimates are the best measure of the underlying behaviour of the series, and the labour market.

Trend estimates are produced by smoothing the seasonally adjusted series using a statistical procedure based on Henderson moving averages. At each survey cycle, the trend estimates are calculated using a centred x-term Henderson moving average of the seasonally adjusted series. The moving averages are centred on the point in time at which the trend is being estimated. The number of terms used to calculate the trend estimates varies across surveys. Generally, ABS monthly surveys use a 13-term Henderson moving average, and quarterly surveys use a 7-term Henderson moving average.

Estimates for the most recent survey cycles cannot be directly calculated using the centred moving average method, as there are insufficient data to do so. Instead, alternative approaches that approximate the smoothing properties of the Henderson moving average are used - such as asymmetric averages. This can lead to revisions in the trend estimates for

the most recent survey cycles, until sufficient data are available to calculate the trend using the centred Henderson moving average. Revisions of trend estimates will also occur with revisions to the original data and re-estimation of seasonal adjustment factors.

---

## Reliability of estimates

The accuracy of an estimate refers to how close that estimate is to the true population value. Where there is a discrepancy between the value of the sample estimate and the true population value, the difference between the two is referred to as the 'error of the sampling estimate'. The total error of the survey estimate results from two types of error:

- sampling error - errors which occur because data were obtained from only a sample rather than the entire population, and
- non-sampling error - errors which occur at any stage of a survey, and can also occur in censuses.

## Sampling error

Sampling error equals the difference between the estimate obtained from a particular sample, and the value that would be obtained if the whole survey population were enumerated. It is important to consider sampling error when publishing survey results as it gives an indication of the accuracy of the estimate, and therefore reflects the importance that can be placed on interpretations. For a given estimator and sample design, the expected size of the sampling error is affected by how similar the units in the target population are and the sample size.

## Variance

Variance is a measure of sampling error that is defined as the average of the squares of the deviation of each possible estimate (based on all possible samples for the same design) from the expected value. It gives an indication of how accurate the survey estimate is likely to be, by measuring the spread of estimates around the expected value. For probability sampling, an estimate of the variance can be calculated from the data values in the particular sample that is generated.

Methods used to calculate estimates of variance in ABS labour-related surveys are outlined below.

- Jack-knife: This method starts by dividing the survey sample into a number of equally sized groups (replicate groups), containing one or more units. Pseudo-estimates of the population total are then calculated from the sample by excluding each replicate group in turn. The jack-knife variance is derived from the variation of the respective pseudo-estimates around the estimate based on the whole sample. This method is used in a

number of household surveys, including the LFS (from November 2002), supplementary surveys (from August 2005), the Multipurpose Household Survey (MPHS) and some labour-related business surveys.

- **Bootstrap:** The Bootstrap is a variance estimation method which relies on the use of replicate samples, essentially sampling from within the main sample. Each of these replicate samples is then used to calculate a replicate estimate and the variation in these replicate estimates is used to calculate the variance of a particular estimate.
- **Ultimate cluster variance:** This method is used in some multi-stage sampling, and involves using the variation in estimates derived from the first-stage units to estimate the variance of the total estimate. This method is used in the Survey of Employee Earnings and Hours.
- **Split halves:** This method involves dividing the sample into half and, from each half, obtaining an independent estimate of the total. The variance estimate is produced using the square of the difference of these estimates. Variations of the split halves method for calculating variance estimates were used in a number of household surveys, including the LFS prior to November 2002 and supplementary surveys prior to August 2005.

The variances indicated in ABS household survey publications are generally based on models of each survey's variance. The variances for a range of estimates are calculated using one of the above methods, and a curve is fitted to the results. This curve indicates the level of variance which could be expected for a particular size of estimate.

## **Standard Error (SE)**

The most commonly used measure of sampling error is called the standard error (SE). The SE is equal to the square root of the variance. An estimate of the SE can be derived from either the population variance (if known) or the estimated variance from the sample units. Any estimate derived from a probability based sample survey has an SE associated with it (called the SE of the estimate). The main features of SEs are set out below.

- SEs indicate how close survey estimates are likely to be to the expected population values that would be obtained from a census conducted under the same procedures and processes;
- SEs provide measures of variation in estimates obtained from all possible samples under a given design;
- Small SEs indicate that variation in estimates from repeated samples is small, and it is likely that sample estimates will be close to the true population values, regardless of the sample selected;
- Estimates of SEs can be obtained from any probability sample - different random samples will produce different estimates of SEs;
- SEs calculated from survey samples are themselves estimates, and thus also subject to SEs;

- When comparing survey estimates, statements should be made about the SEs of those estimates; and
- SEs can be used to work out confidence intervals. This concept is explained below.

### **Confidence Interval (CI)**

A confidence interval (CI) is defined as an interval, centred on the estimate, with a prescribed level of probability that it includes the true population value (if the estimator is unbiased), or the mean of the sampling distribution (if the estimator is biased). Estimates from ABS surveys are usually unbiased.

Estimates are often presented in terms of a CI. Most commonly, CIs are constructed for 66%, 95%, and 99% levels of probability. The true value is said to have a given probability of lying within the constructed interval. For example:

- 66% chance that the true value lies within 1 standard error of the estimate (2 chances in 3);
- 95% chance that the true value lies within 2 standard errors of the estimate (19 chances in 20); and
- 99% chance that the true value lies within 3 standard errors of the estimate (99 chances in 100).

CIs are constructed using the standard error associated with an estimate. For example, a 95% CI is equivalent to the survey estimate, plus or minus two times the standard error of the estimate. For example, the originally published LFS estimate of employment (seasonally adjusted) for September 2017 was 12,290,200 persons, and the estimate had a standard error of 44,400. The 95% CI could be expressed: "we are 95% confident that the true value for employment lies between 12,201,400 and 12,379,000".

### **Relative Standard Error (RSE)**

Another measure of sampling error is the relative standard error (RSE). This is the standard error expressed as a percentage of the estimate. Since the standard error of an estimate is generally related to the size of the estimate, it is not possible to deduce the accuracy of the estimate from the standard error without also referring to the size of the estimate. The relative standard error avoids the need to refer to the estimate, since the standard error is expressed as a proportion of the estimate. RSEs are useful when comparing the variability of population estimates of different sizes. They are commonly expressed as percentages.

Very small estimates are subject to high RSEs, which detract from their usefulness. In some ABS labour-related statistical publications, estimates with an RSE greater than 25% but less than 50% have an asterisk (\*) displayed beside the estimate, indicating they should be used with caution. Estimates with an RSE greater than 50% have two asterisks (\*\*) displayed

beside the estimate, indicating they are so unreliable as to detract seriously from their value for most reasonable uses. All cells in a Data Cube with RSEs greater than 25% contain a comment indicating the size of the RSE. These cells are identified by a red indicator in the corner of the cell. The comment appears when the mouse pointer hovers over the cell.

## **Non-sampling error**

Non-sampling error refers to all other errors in the estimate. Non-sampling error can be caused by non-response, badly designed questionnaires, respondent bias, interviewer bias, collection bias, frame deficiencies and processing errors. It is often difficult and expensive to quantify non-sampling error.

Non-sampling errors can occur at any stage of the process, and in both censuses and sample surveys. Non-sampling errors can be grouped into two main types: systematic and variable. Systematic error (called bias) makes survey results unrepresentative of the population value by systematically distorting the survey estimates. Variable error can distort the results on any given occasion, but tends to balance out on average over time.

Every effort is made to minimise non-sampling error in ABS surveys at every stage of the survey, through careful design of collections, and the use of rigorous editing and quality control procedures in the compilation of data. Some of the approaches adopted are listed below.

- Reducing frame deficiencies.
- Reducing non-response - Non-response results in bias in the estimate because it is possible the non-respondents have different characteristics to respondents, leading to an under-representation of the characteristics of non-respondents in the sample survey estimate. The ABS pursues a policy of intensive follow up of non-respondents. This includes multiple visits or telephone calls in an attempt to contact respondents, and letters requesting compliance with the survey. Partial non-response is also followed up with respondents.
- Reducing instrument errors - These errors relate to poor questionnaire design, leading to questions which are not easily understood by respondents, and hence incorrect responses. This is particularly relevant for household surveys. The ABS ensures that all household survey questionnaires are carefully tested using cognitive testing and dress rehearsals of the survey before it is officially conducted. New business survey questionnaires and additional questions in business surveys are also rigorously tested before they are introduced.

## **Measures of non-sampling error**

Non-sampling error is difficult to quantify; however, an indication of the level of non-sampling error can be determined from a number of quality measures. These include:

- Response rates: The number of responding units in a survey expressed as a proportion of the total number of units selected (excluding deaths). Response rates can also be calculated for individual questions within a survey.
- Imputation rates: The number of responses which need to be imputed expressed as a proportion of the total number of responses
- Coverage rates: An estimate of the proportion of units in the target population which are not covered by the frame
- Any Responsible Adult rates: The number of responding units in a survey for which information was supplied by a responsible adult rather than personally, expressed as a proportion of the total number of responding units. Any Responsible Adult rates can only be calculated for household surveys.

---

## Confidentiality

All releases of data from the ABS are confidentialised to ensure that no unit (e.g. person or business) is able to be identified. The ABS applies a set of rules, concerning the minimum number of responses required to contribute to each data cell of a table, and the maximum proportion that any one respondent can contribute to a table cell, to ensure that information about specific units cannot be derived from published survey results.

In some instances it is not possible to confidentialise responses from businesses that contribute substantially to a data cell. In this case, agreement is sought from the business for their data to still be published. If agreement is not reached, all affected data cells are suppressed.

Under the Census and Statistics Act, 1905 it is an offence to release any information collected under the Act that is likely to enable identification of any particular individual or organisation. Introduced random error is used to ensure that no data are released which could risk the identification of individuals in the statistics.

A technique, known as perturbation, has been developed to randomly adjust cell values. Random adjustment of the data is considered to be the most satisfactory technique for avoiding the release of identifiable data. When the technique is applied, all cells are slightly adjusted to prevent any identifiable data being exposed. These adjustments result in small introduced random errors. However, the information value of the table as a whole is not impaired.

These adjustments may cause the sum of rows or columns to differ by small amounts from table totals. The counts are adjusted independently in a controlled manner, so the same information is adjusted by the same amount. However, tables at higher geographic levels may not be equal to the sum of the tables for the component geographic units.



It is not possible to determine which individual figures have been affected by random error adjustments, but the small variance which may be associated with derived totals can, for the most part, be ignored.